arXiv: submitted to arXiv.org [cs.CV] 18 Dec 2012

Visual Objects Classification with Sliding Spatial Pyramid Matching

Hao Wooi Lim[1] and Yong Haur Tay[2]
[1] haowooilim@1utar.my, [2] tayyh@utar.edu.my
Computer Vision & Intelligent Systems (CVIS) group, Faculty of Engineering and Science, University of Tunku Abdul Rahman

Abstract

We present a method for visual object classification using only a single feature, transformed color SIFT [15] with a variant of Spatial Pyramid Matching (SPM) that we called Sliding Spatial Pyramid Matching (SSPM), trained with an ensemble of linear regression (provided by LINEAR) to obtained state of the art result on Caltech-101 [22] of 83.46%. SSPM is a special version of SPM where instead of dividing an image into K number of regions, a subwindow of fixed size is slide around the image with a fixed step size. For each subwindow, a histogram of visual words is generated. To obtained the visual vocabulary, instead of performing K-means clustering [26], we randomly pick N exemplars from the training set and encode them with a soft non-linear mapping method from [16]. We then trained 15 models, each with a different visual word size with linear regression. All 15 models are then averaged together to form a single strong model.

1 Introduction

Recently, there has been a lot of development in the area of Deep Learning (DL). Not unlike traditional Multi-layer Neural Networks, DL sought to solve the difficult problems such as visual objects classification by utilizing many layers of neurons, which has been traditionally too computationally intensive to train, with the help of fast, many-cores Graphics Processing Unit (GPU) to help speedup training. One interesting outcome of the research in DL is such that, DL allows for learning of feature detectors, thus mostly eliminating the need for hand-engineered features. And sure enough, recent papers such as [18] has shown that DL could even obtained state of the art results on various visual object data set such as MNIST [19], NIST SD 19 [20], CIFAR-10 [21] and etc. However, noticed that most of the data set tested are usually consists of very small images. This is due to the fact that data set with huge dimensionality would pose an even bigger challenge for GPU, sometimes even requiring weeks of training, even with a cluster of GPUs. Moreover, not everyone has a huge cluster of GPUs at their disposal.

Looking at some of the best-performing methods ([1], [4], [8], [9]), shown on Fig. 1 on Caltech-101 [22], one trend seems apparent: many different features are usually employed to obtained the published results, most notably Scale-invariant feature transform (SIFT) [23], Pyramid histogram of orientated gradients (PHOG) [9], Self-similarity (SSIM) [25], Geometric Blur (GB) [24] and etc. While it is easy to obtain better accuracy by employing many different features, it significantly increases the computation power required to train them.

arXiv: some other text goes here

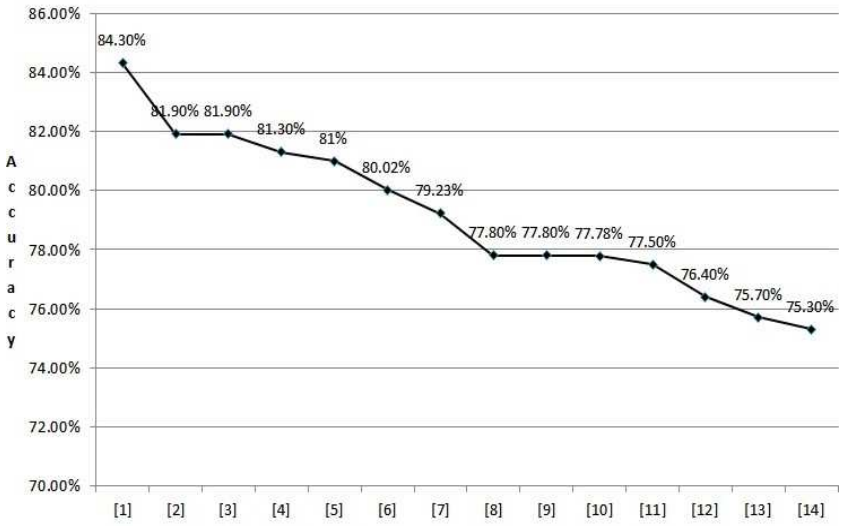

Figure 1: Figure compares the accuracy on Caltech-101 data set, all trained on 30 samples.

Some algorithm like kernel-based Support Vector Machine (SVM) which requires $O(nd)$ testing time, where n is the number of support vectors and d is the dimensionality of the feature vector [27], it is just not feasible when the model is complex, resulting in large number of support vectors or when the feature dimension is large. Unfortunately, it is pretty easy to feed a lot of features for training, allowing the training process to discover salient features useful for accurate classification. Thus, in most cases, it is difficult to obtain good results with kernel-based SVM without also making training intractable.

In this paper, our goal is to lower the error bound siginificantly on challenging visual object data set such as Caltech-101 [22] that has high enough number of categories and dimensionality, while also employing methods that are much more efficient to train than DL-based methods, relying on hand-engineered feature but refrained from combining many different kind of features so that training remains quick and efficient.

2 Our method

Here, we describe our method, along with the best parameters that we have found.

Our method begins with extraction of features. Following the method employed by [15], we use Harris-Laplace salient point detector to localize all the local keypoints in each image. For each keypoint, we use the Transformed Color SIFT as the descriptor.

To obtained a visual vocabulary, we randomly pick N exemplars from the training set instead of performing K-means clustering [26], which was the usual method employed by most paper.

To encode keypoint into feature vector, we encode them with a soft non-linear mapping method as described in [16].

Next, we use a special version of Spatial Pyramid Matching (SPM) where instead of dividing an image into K number of regions whereby for each region, a histogram of visual words is generated, a subwindow of fixed size is slide around the image with fixed step size.

arXiv: some other text goes here

Each subwindow then, is process the same as you would with a region in SPM. We call this Sliding Spatial Pyramid Matching (SSPM). Here, we set the width and height of the subwindow to be one-third of the original image's width and height respectively, while step size x and y to be the one-third of the subwindow's width and height respectively.

Since for each region we obtained a feature vector with size equal to the visual vocabulary size, we append each feature vector from each region into a single feature vector to be trained by linear regression.

To train the system, we used LIBLINEAR [17] for this purpose (Due to the high dimensionality of our feature vector, kernel-based SVM is unfeasible). We choose L2-regularized logistic regression (primal) with probabilistic output.

Finally, to improve the accuracy, we trained multiple models and combined them by multiplying together the probabilistic output to obtained the final probability distribution over class label.

3 Data sets and Experimental Protocol

We conduct our experiment on Caltech-101 [22]. The data set consists of images from 101 object categories, and contains from 31 to 800 images per category. Most images are medium resolution, about 300×300 pixels. The significance of this database is its large inter-class variability. Following standard procedures, the Caltech-101 data is split into 30 training images (chosen randomly) per category and 50 for testing. BACKROUND_GOOGLE category was not included in the experiment. The classification process is repeated 4 times (changing the training and test sets) and the average performance score is reported.

4 Results

We study the influence of a variable visual vocabulary size ($N = \{1000, 900, ..., 100\}$) or fixed visual vocabulary size ($N = 1000$) on several models that was then averaged together to obtain a final result. Note that each model has a different visual vocabulary due to the fact that visual word is obtained with randomly chosen exemplars. To speedup experiment, SIFT is used in this case.

As shown in Tab. 1, variable visual word seems to performed slightly better. This makes sense, when combining several models together, models that are different from each other tends to performed better after being combined.

We then experiment with different step size and subwindow size. For method A, we set the width and height of the subwindow to be half of the original image's width and height respectively, while step size x and y to be the half of the subwindow's width and height respectively. For method B, we set the width and height of the subwindow to be one-third of the original image's width and height respectively, while step size x and y to be the one-third of the subwindow's width and height respectively. As one might exect, method B would result in features with higher dimension compared to method A, due to the fact that method B generates more regions. To speedup experiment, SIFT is used in this case.

As shown in Tab. 2, method B seems to be performing better, with almost negligible increase in training time due to the higher dimension feature vector generated. Interestingly,

arXiv: some other text goes here

Number of models	Fixed visual vocabulary size	Variable visual vocabulary size
1	71.13%	70.92%
2	73.18%	72.95%
3	73.37%	72.70%
4	73.67%	72.97%
5	73.61%	73.69%
6	73.65%	73.57%
7	73.67%	73.54%
8	73.80%	73.71%
9	73.69%	73.48%
10	73.78%	**73.84%**

Table 1: Table compares the accuracy using variable visual vocabulary size vs fixed visual vocabulary size.

Number of models	Method A	Method B
1	70.92%	72.78%
2	72.95%	73.95%
3	72.70%	74.65%
4	72.97%	74.90%
5	73.69%	75.03%
6	73.57%	75.45%
7	73.54%	75.30%
8	73.71%	**75.64%**
9	73.48%	75.41%
10	73.84%	75.47%

Table 2: Table compares the accuracy using smaller subwindow size vs larger subwindow size.

it also seems to show there is diminishing returns with averaging more models. Once more than 8 models are averaged together, accuracy drops slightly.

Next, we determine the best performing local keypoint descriptors. We test SIFT, HSV-SIFT, Opponent-SIFT and Transformed Color SIFT.

As shown in Tab. 3, color variants of SIFT seems to be performed significantly better than SIFT, indicating that color information is a pretty strong discriminating factor.

Finally, we investigate if combining more models would improve the accuracy even further. In method C, we combine 10 models and set visual vocabulary size to Eq. 1, while in method D, we combine 15 models and set visual vocabulary size to Eq. 2. In both cases, we used Transformed Color Sift as the local keypoint descriptor.

$$N = 1000, 1000, 800, 800, ..., 200, 200 \tag{1}$$

$$N = 1000, 1000, 1000, 800, 800, 800, ..., 200, 200, 200 \tag{2}$$

arXiv: some other text goes here

Number of models	SIFT	HSV-SIFT	Opponent-SIFT	Transformed color SIFT
1	72.27%	74.22%	78.18%	79.03%
2	73.97%	76.64%	79.7%	81.38%
3	74.90%	77.44%	80.72%	82.44%
4	75.20%	77.89%	81.1%	82.93%
5	75.32%	77.91%	81.70%	82.99%
6	75.15%	78.01%	82.10%	82.93%
7	75.49%	78.20%	81.9%	82.76%
8	75.66%	78.88%	81.99%	82.95%
9	75.62%	78.86%	82.27%	83.12%
10	75.85%	78.94%	82.44%	**83.14%**

Table 3: Table compares the accuracy using different local keypoint descriptors.

Method C	Method D
83.13%	**83.46%**

Table 4: Table compares the accuracy using different number of models.

As shown in Tab. 4, perhaps unsurprisingly, combining more models did help to improve the accuracy. We tried increasing the number more, unfortunately, we hit a point of diminishing returns.

On a final note, the reason having more models help obtained a higher accuracy, is most probably due to the fact that since visual words are obtained randomly, each model's feature was encoded with a different set of visual word, thereby giving each model a different viewpoint of the problem.

In terms of training time, the test program which was coded in C++ and speed up with OpenMP to take advantage of all 4 cores, took about 5 hours to train for the 10-model classifier and about 8 hours for the 15-model classifier on a single quad-core Intel Core i5-2500K machine.

5 Conclusion & Future works

To conclude, we have presented our proposed visual object classification method that utilizes only a single best-performing local keypoint descriptor (Transformed color SIFT) while being efficient to train. We managed to obtained state of the art result on Caltech-101 of 83.46%, showing the potential of such architecture. We suspect better results might be possible if multiple models utilizes different type local keypoint descriptor. However, it is unclear yet if significantly better results would be possible or if it would slow down training too much.

References

[1] J. Yang, Y. Li., Y. Tian, L. Duan and W. Gao. Group sensitive multiple kernel learning for object categorization. In ICCV, 2009

arXiv: some other text goes here

[2] S. Todorovic and N. Ahuja. Learning subcategory relevances for category recognition. In CVPR, 2008

[3] F. Li, J. Carreira and C. Sminchisescu. Object Recognition as Ranking Holistic Figure-ground Hypotheses. In CVPR, 2010

[4] A. Bosch, A. Zisserman and X. Munoz. Image classification using random forests and ferns. In ICCV, 2007

[5] S. McCann and D. G. Lowe. Spatially Local Coding for Object Recognition. In ACCV, 2012

[6] Gabriel L. Oliveira, Erickson R. Nascimento, Antonio W. Vieira and Mario F. M. Campos. Sparse spatial coding: A novel approach for efficient and accurate object recognition. In ICRA, 2012

[7] O. Boiman, E. Shechtman and M. Irani. In defense of nearest-neighbor based image classification. In CVPR, 2008

[8] P. Gehler and S. Nowozin. On Feature Combination for Multiclass Object Classification. In ICCV, 2009

[9] A. Bosch, A. Zisserman and X. Munoz. Representing shape with a spatial pyramid kernel. In CIVR, 2007

[10] K. Chatfield, V. Lempitsky, A. Vedaldi and A. Zisserman. The devil is in the details: an evaluation of recent feature encoding method. In BMVC, 2011

[11] L. Bo, K. Lai, X. Ren and D. Fox. Object recognition with hierarchical kernel descriptors. In CVPR, 2011

[12] L. Bo, X. Ren and D. Fox. Kernel Descriptors for Visual Recognition. In NIPS, 2010

[13] Y. Boureau, F. Bach, Y. LeCun and J. Ponce. Learning mid-level features for recognition. In CVPR, 2010

[14] T. Jia, C. Huang and T. Darrell. Beyond spatial pyramids: Receptive field learning for pooled image features. In CVPR, 2012

[15] Koen E. A. van de Sande, Theo Gevers and Cees G. M. Snoek. Evaluating Color Descriptors for Object and Scene Recognition. In PAMI, 2010

[16] Adam Coates, Honglak Lee and Andrew Ng. An analysis of single-layer networks in unsupervised feature learning. In NIPS, 2010

[17] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang and C.-J. Lin. LIBLINEAR: A library for large linear classification. In JMLR, 2008.

[18] D. C. Ciresan, U. Meier and J. Schmidhuber. Multi-column deep neural networks for image classification. In CVPR, 2012

arXiv: some other text goes here

[19] Y. LeCun, L. Bottou, Y. Bengio and P. Haffner. Gradient-based learning applied to document recognition. In Proceedings of the IEEE, 86(11):22782324, November 1998

[20] P. J. Grother. NIST special database 19 - Handprinted forms and characters database. Technical report, National Institute of Standards and Thechnology (NIST), 1995

[21] A. Krizhevsky. Learning multiple layers of features from tiny images. Masters thesis, Computer Science Department, University of Toronto, 2009

[22] L. Fei-Fei, R. Fergus and P. Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In IEEE CVPR Workshop of Generative Model Based Vision, 2004

[23] D. G. Lowe. Object recognition from local scale-invariant features. In ICCV, 1999

[24] A. C. Berg, T. L. Berg and J. Malik. Shape matching and object recognition using low distortion correspondences. In CVPR, 2005

[25] E. Shechtman and M. Irani. Matching local self-similarities across images and videos. In CVPR, 2007

[26] K. Alsabti, S. Ranka and V. Singh. An efficient k-means clustering algorithm. In Proceedings of the First Workshop on High Performance Data Mining, Orlando, FL, 1998

[27] H. Cao, T. Naito and Y. Ninomiya. Approximate RBF kernel SVM and its applications in pedestrian classification. In 1st International Workshop on Machine Learning for Vision-based Motion Analysis, 2008.